

APPLICATION
FOR
UNITED STATES LETTERS PATENT

**TITLE: A METHOD AND APPARATUS FOR
 SPECULATIVE TRACING**

APPLICANTS: Bryan M. CANTRILL

32615
PATENT TRADEMARK OFFICE

“EXPRESS MAIL” Mailing Label Number: EL974017475US

Date of Deposit: November 14, 2003

A METHOD AND APPARATUS FOR SPECULATIVE TRACING

Background

[0001] A tracing framework is used to understand the behavior of complex computer systems. Tracing performed using the tracing framework involves recording data from a location or occurrence in the software of the computer system. In a tracing framework that offers comprehensive coverage, the framework provides a mechanism to allow events *not* to be traced or the user may be flooded with unwanted data.

[0002] Tracing frameworks use predicates to limit the events traced by only tracing data if a certain condition is found to be true. Predicates are a useful mechanism when the user knows whether the probe event is interesting when a probe is triggered. For example, if the user is only interested in an activity associated with a certain process or a certain file descriptor, the user can define a predicate that obtains knowledge about only that particular process or file descriptor when the probe is triggered.

[0003] However, some situations exist when the user may not have knowledge whether a given probe event is interesting at the time that the probe is triggered. Often the user can only make the determination that the probe event is interesting after the probe is triggered.

[0004] For example, if a function with the software of a computer system is failing in an irregular manner with a common error code, the user may wish to better understand the code path that is leading to the error condition. To capture the code path, the user could enable every probe—but only if the failing function call can be isolated in such a way that meaningful predicates can be defined. If the failures are sporadic or non-deterministic, the user is forced to trace all events that may be

interesting, and later use post-process analysis of the data to filter out particular data that is not associated with the failing code path. In this case, even though the number of interesting events may be reasonably small, the number of events to analyze during post-process analysis and that must be traced is very large.

Summary

[0005] In general, one aspect of the invention relates to a method of a speculative tracing. The method includes defining the speculative tracing using a plurality of probes, firing at least one of the plurality of probes defined by the speculative tracing, allocating at least one instance of a first speculative buffer arranged to transfer data to a first principal buffer, if one of the plurality of probes comprises a first speculation function, and determining a first state value associated with the first speculative buffer.

[0006] In general, one aspect of the invention relates to a system for a speculative tracing using a tracing framework. The system including a first principal buffer configured to store data from the tracing framework and associated with a first processor, a first instance of a first speculative buffer associated with the first principal buffer and configured to transfer data to the first principal buffer, where the first speculative buffer has a first state value associated therewith, and a plurality of probes defining the speculative tracing executing on the tracing framework, where the first state value is updated upon firing at least one of the plurality of probes.

[0007] In general, one aspect of the invention relates to a computer system for a speculative tracing. The computer system includes a first processor, a memory, a storage device, and software instructions stored in the memory for enabling the computer system to define the speculative tracing using a plurality of probes, fire at least one of the plurality of probes defined by the speculative tracing, allocate at

least one instance of a first speculative buffer arranged to transfer data to a first principal buffer, if one of the plurality of probes comprises a first speculation function, and determine a first state value associated with the first speculative buffer.

[0008] Other aspects of the invention will be apparent from the following description and the appended claims.

Brief Description of Drawings

[0009] Figure 1 shows a tracing framework in accordance with one or more embodiments of the invention.

[0010] Figure 2 shows a principal buffer and a speculative buffer of a tracing framework in accordance with one or more embodiments of the invention.

[0011] Figures 3 shows speculative buffers associated with a single principal buffer in accordance with one or more embodiments of the invention.

[0012] Figure 4 shows a state transition diagram of a speculative tracing operation for a multi-processor environment in accordance with one or more embodiments of the invention.

[0013] Figure 5 shows a networked computer system in accordance with one or more embodiments of the invention.

Detailed Description

[0014] Specific embodiments of the invention will now be described in detail with reference to the accompanying figures. Like elements in the various figures are denoted by like reference numerals for consistency.

[0015] In the following detailed description of embodiments of the invention, numerous specific details are set forth in order to provide a more thorough

understanding of the invention. However, it will be apparent to one of ordinary skill in the art that the invention may be practiced without these specific details. In other instances, well-known features have not been described in detail to avoid obscuring the invention.

[0016] One or more embodiments of the invention relate to a tracing framework and a method for tracing operation. Specifically, one or more embodiments of the invention relate to a tracing framework and method for speculative tracing.

[0017] Figure 1 shows a tracing framework in accordance with one or more embodiments of the invention. Figure 1 shows source code (100), which defines performance-related questions with respect to a software system. Performance-related questions, for example, may relate to processing speed, resource consumption, and/or proper execution, *etc.*

[0018] In an embodiment of the invention, the source code (100) is obtained using a command line or a graphical user interface. Once obtained, the source code (100) is compiled into executable object code (102). Requests from the object code (102) are communicated to an execution-level tracing framework (108) via an execution interface (106). The tracing framework (108) interprets the requests from the object code (102) and forwards the requests to the probe providers (110), which activate certain probes (112 and 114) in an instrumented program (not shown). The probes (112 and 114) correspond to a particular location and/or activity within the instrumented program, and answer the specific performance-related question. Further, these probes (112 and 114) can gather the specified information and store the information accordingly.

[0019] In one or more embodiments, the probes (112 and 114) may be described in the following psuedo-code:

Code Sample 1

1 probe description

```
2  /predicate/  
   {  
3  action  
   }
```

[0020] In line 1 of the Code Sample 1, the name of the probe is defined. Line 2 shows a predicate, which is a pre-determined, conditional statement that determines whether the action (shown at line 3) of the probe is to be executed. The predicate evaluates to a value of true or false, *e.g.*, an integer value of zero or one, or a defined pointer type. Line 3 of Code Sample 1 defines the action of the probe that is executed, *i.e.*, the tracing operation. Examples of tracing operations include tracing (or recording data), modifying a state variable, *etc.*

[0021] A tracing operation is considered to be initiated once the predicate is evaluated. A tracing operation may be terminated if a pre-determined condition, as defined by the predicate, is not satisfied. In this manner, only relevant data is traced and subsequently stored. However, if a predicate is not necessarily useful in determining whether the data is relevant or desirable, then the probes are triggered. As a result, the corresponding data is traced, stored, and evaluated.

[0022] In one or more embodiments, the invention allows relevant and/or desirable data to be stored after the probe has been fully executed, using a mechanism known as speculative tracing. Speculative tracing allows data to be “tentatively” traced and then later committed or discarded based on whether the data is desirable or not. Therefore, probes in the invention can be classified as either speculative or non-speculative in nature.

[0023] The invention uses specialized storage areas to execute speculative tracing. Figure 2 shows a principal buffer and a speculative buffer of a tracing framework in accordance with one or more embodiments of the invention. In particular, Figure 2 shows a principal buffer (200), which is associated with a particular processor (not shown), and a speculative buffer (202), which is associated with the

principal buffer (200). Both the principal buffer (200) and the speculative buffer (202) are connected to a tracing framework (108), which has functionality to control the buffers and transfer data to and between the buffers.

[0024] In one or more embodiments of the invention, the arrangement of the buffers and the association between the buffers and processors may vary. For example, a buffer may be associated with multiple processors, or multiple buffers may be associated with a single processor. Further, in one or more embodiments, multiple speculative buffers may be stored in cells of the same array or table.

[0025] The speculative buffer (202) stores data generated during speculative tracing. Generally, the speculative buffer (202) only stores data for a temporary time period. On the other hand, the principal buffer (200) stores data during tracing, whether speculative or non-speculative. In one embodiment of the invention, upon satisfying a user's pre-defined condition for transfer, data is transferred from the speculative buffer (202) to the principal buffer (200).

[0026] Additionally, in one or more embodiments, the tracing framework (108) may also include a drop counter associated with the principal buffer (200). The drop counter may increment when no space is available for data to be transferred from the speculative buffer (202) to the principal buffer (300). As a result, no data is transferred to the principal buffer.

[0027] Additionally, in one or more embodiments, the speculative buffer (202) is a finite resource of the tracing framework. Thus, a system error may occur (*i.e.*, an error counter is incremented) when no speculative buffer (202) is available during an attempt to allocate the speculative buffer (202).

[0028] In one or more embodiments, speculative tracing may be described with the following psuedo-code:

Code Sample 2

```
1      probe description 1
      {
2      buffer name = speculation();
      }

3      probe description 2
4      speculate (buffer name);
      {
5      action;
      }

6      probe description 3
7      /conditional 1/
      {
8      commit (buffer name);
      }

9      probe description 4
10     /conditional 2/
      {
11     discard(buffer name);
      }
```

[0029] In the above psuedo-code, speculative tracing is defined by a plurality of probe descriptions, which may include particular functions, *i.e.*, *speculation ()*, *speculate ()*, *commit ()*, and *discard ()*.

[0030] As shown in lines 1 and 2 of Code Sample 2 (*i.e.*, probe description 1), a *speculation()* function is called, which allocates a speculative buffer and returns an identifier for the speculative buffer. The *speculate()* function as shown in probe description 2 at lines 3-5, may be triggered, if and only if, an identifier from the *speculation()* function is passed to the *speculate()* function. Further, the *speculate()* function denotes that the remainder of the clause should be traced to the speculative buffer specified by the *speculation()* function.

[0031] In one or more embodiments, speculative tracing may define more than one *speculate()* function. Further, if more than one probe having a *speculate()* function fires, the data (as defined by the function) is traced to a speculative buffer

(having the same identifier), which is associated with the processor on which the probe fired (assuming the same identifier is passed to the *speculate()* function).

[0032] Lines 6-8 of Code Sample 2 show probe description 3 having the *commit()* function. The *commit()* function commits the speculative buffer associated with the returned identifier to the principal buffer, depending on a conditional statement. In other words, the data stored in the speculative buffer is transferred to its corresponding principal buffer.

[0033] Lines 9-11 of Code Sample 2 show probe description 4 having the *discard()* function. The *discard()* function discards the data in the speculative buffer associated with the returned identifier, depending on a conditional statement. In other words, the data in the speculative buffer is discarded and waits to be “cleaned.”

[0034] One of ordinary skill in the art will appreciate that the probe descriptions may contain predicates. Additionally, one of ordinary skill in the art will understand that an identifier from a *speculation()* function must be passed to a *speculate()* function and a *speculate()* function may not follow any data recording actions.

[0035] Further, one skilled in the art will appreciate that more than one *speculation()* function may be called thereby returning the proper identifiers for the respective speculative buffers. In other words, multiple *speculation()* functions may be called allocating uniquely identified speculative buffers, which may be associated with the same principal buffer.

[0036] For example, Figure 3 shows two speculative buffers associated with a single principal buffer in accordance with one or more embodiments of the invention. The tracing framework (108) is connected to a principal buffer (200) and two speculative buffers (202 and 204). The speculative buffers (202 and 204) are each connected to the principal buffer (200), which is associated with a single

processor (not shown). Further, the tracing framework (108) includes functionality to transfer data from speculative buffers (202 and 204) to the principal buffer (200) associated with the single processor (not shown). Again, the arrangement of the buffers and the association between the buffers and processors may vary, in accordance with one or more embodiments of the invention.

[0037] Additionally, each speculation has a global state, which updates during speculative tracing and governs the execution of the speculative tracing. Table 1 show a chart of global states of a speculation in accordance with one or more embodiments of the invention. In one or more embodiments of the invention, the state names may vary and be grouped in a manner to be represented by a single global state.

Table 1- Global State of Speculation

Global State	Description
Inactive	Initial state; inactive speculative tracing operation
Active	Speculative buffer allocated, however, no tracing has occurred
Speculate-One	Speculative tracing operation has occurred on one CPU
Speculate-Many	Speculative tracing operation has occurred on multiple CPUs
Commit-One	Results of speculative tracing operation being committed to principal buffer associated with one CPU
Commit-Many	Results of speculative tracing operation being committed to principal buffer associated with multiple CPUs
Discard	Results of speculative tracing operation being discarded

[0038] Figure 4 shows a state diagram of global states listed in Table 1, in accordance with one or more embodiments of the invention. Figure 4 begins in the “inactive” state (400), *i.e.*, no operation or non-speculative tracing is being executed by any processor. Once a probe having a *speculation()* function is

triggered, a speculative buffer is allocated, an identifier for the speculative buffer is returned, and the state updates to “active” (402).

[0039] From this point, a probe having a *commit()* function or a *discard()* function may be encountered. The state may be updated to “commit-one” state (412) or “discard” state (410), respectively. Based on the state, the data, if any, is committed or discarded, accordingly.

[0040] Alternatively, from the “active” state (402), a “speculate-one” state (404) may be reached when a probe having the *speculate()* function is triggered. As previously mentioned, the *speculate()* function results in data being traced and stored in the speculative buffer. Transitioning into the “speculate-one” state records the processor on which the *speculate()* function occurred. From the “speculate-one” state (404), a transition to a “commit-one” state (412) or a “discard” state (410) may occur.

[0041] However, if another probe description having a *speculate()* function is encountered, *i.e.*, if another probe speculates to the same buffer on a different processor, then (and only then) is the “speculate-one” state (404) transitioned to the “speculate-many” state (406). Similar to the “speculate-one” state (404), the “speculate-many” state (406) may transition to a “commit-many” state (408) or a “discard” state (410). In the “commit-many” state (408), the data in the speculative buffers is transferred to their respective principal buffers. Alternatively, from the “discard” state (410), the data in the speculative buffers are discarded, and, subsequently, cleaned.

[0042] In one or more embodiments, the global states optimize for the case that a speculative buffer is only active on one processor at the time of a *commit()* function or a *discard()* function. In this case, other processors need not take any action, and the *speculation()* function is immediately ready for reuse.

[0043] On the other hand, in one or more embodiments, speculative tracing may occur where the speculative buffer is used on one more than one processor. In this case, the speculative buffers associated with the *speculation()* function must be asynchronously cleaned (which may potentially lead to a higher rate of dirty speculative drops). The rate at which the speculative buffers are cleaned occurs at a user-configurable, fixed interval (*i.e.*, not at probe-triggering time) by making a call to each processor to atomically reset each speculative buffer. Once a speculative buffer has been committed or discarded, the speculative buffer cannot be reused until all of the processors take the same action on their respective speculative buffers. Accordingly, subsequent *speculation()* function calls will be “silently” discarded and *commit()* and *discard()* function calls will fail thereby incrementing a counter, whose contents may be reported back to the user.

[0044] In one or more embodiments, the data on the processor calling the *commit()* function is transferred immediately, while the data on other processors are transferred some time after the *commit()* function was called. However, this time is guaranteed to be no longer than the time dictated by the cleaning rate.

[0045] In one or more embodiments, the data on the processor calling the *discard()* function is discarded immediately, while the data on the other processors are discarded some time after the *discard()* function was called. However, this time is guaranteed to be no longer than the time dictated by the cleaning rate.

[0046] One skilled in the art will appreciate that the invention may be implemented on virtually any type of computer regardless of the platform being used. For example, as shown in Figure 5, a typical networked computer system (500) includes a processor (502), associated memory (504), a storage device (506), and numerous other elements and functionalities typical of today's computers (not shown). The networked computer (500) may also include input means, such as a keyboard (508) and a mouse (510), and output means, such as a monitor (512).

The networked computer system (500) is connected to a local area network (LAN) or a wide area network (514) (*e.g.*, the Internet) via a network interface connection (not shown). Those skilled in the art will appreciate that these input and output means may take other forms. Those skilled in the art will appreciate that one or more elements of the aforementioned computer (500) may be located at a remote location and connected to the other elements over a network.

[0047] The invention provides a mechanism for a system to be speculatively traced. Speculative tracing also reduces storage of unnecessary data in primary buffers and minimizes analysis time of tracing data. The invention also allows a system to be speculatively traced while reducing lost or dropped data stored in the buffers without unnecessarily locking buffers.

[0048] While the invention has been described with respect to a limited number of embodiments, those skilled in the art, having benefit of this disclosure, will appreciate that other embodiments can be devised which do not depart from the scope of the invention as disclosed herein. Accordingly, the scope of the invention should be limited only by the attached claims.